

On the Performance of Short Block Codes over Finite-State Channels in the Rare-Transition Regime

Fatemeh Hamidi-Sepehr, Jean-Francois Chamberland, Henry D. Pfister

Abstract

As the mobile application landscape expands, wireless networks are tasked with supporting different connection profiles, including real-time traffic and delay-sensitive communications. Among many ensuing engineering challenges is the need to better understand the fundamental limits of forward error correction in non-asymptotic regimes. This article characterizes the performance of random block codes over finite-state channels and evaluates their queueing performance under maximum-likelihood decoding. In particular, classical results from information theory are revisited in the context of channels with rare transitions, and bounds on the probabilities of decoding failure are derived for random codes. This creates an analysis framework where channel dependencies within and across codewords are preserved. Such results are subsequently integrated into a queueing problem formulation. For instance, it is shown that, for random coding on the Gilbert-Elliott channel, the performance analysis based on upper bounds on error probability provides very good estimates of system performance and optimum code parameters. Overall, this study offers new insights about the impact of channel correlation on the performance of delay-aware, point-to-point communication links. It also provides novel guidelines on how to select code rates and block lengths for real-time traffic over wireless communication infrastructures.

This material is based upon work supported by the National Science Foundation (NSF) under Grants No. 0747363 and No. 0830696. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation. This paper was presented in part at the 50th Annual Allerton Conference on Communication, Control, and Computing and the 46th Annual Conference on Information Sciences and Systems.

The authors are with the Department of Electrical and Computer Engineering, Texas A&M University, College Station, TX 77843, USA (emails: f_hamidisepehr@tamu.edu; chmbrlnd@tamu.edu; hpfister@tamu.edu).

I. INTRODUCTION

With the ever increasing popularity of advanced mobile devices such as smartphones and tablet personal computers, the demand for ubiquitous low-latency, high-throughput wireless services is growing rapidly. The shared desire for a heightened user experience, which includes real-time applications and mobile interactive sessions, acts as a motivation for the study of highly efficient communication schemes subject to stringent delay constraints. An important aspect of delay-sensitive traffic stems from the fact that intrinsic delivery requirements may preclude the use of very long codewords. As such, the insights offered by classical information theory and based on Shannon capacity are of limited value in this context. A prime goal of this article is to develop a better understanding of delay-constrained communication, queue-based performance criteria and service dependencies attributable to channel memory. In particular, we seek to derive meaningful performance limits for delay-aware systems operating over channels with memory. The emphasis is put on identifying upper bounds on the probabilities of decoding failure for systems employing short block lengths. This is an essential intermediate step in the characterization of queueing behavior for contemporary communication systems.

Computing the probability of decoding failure for specific communication channels and particular coding schemes is of great interest. This line of work dates back to the early days of information theory [1] and has received significant attention in the past, with complete solutions in some cases. One approach that has been remarkably successful, consists of deriving exponential error bounds on the behavior of asymptotically long codewords. This approach was popularized by Gallager [2] and these bounding techniques have been applied to both memoryless channels and finite-state channels with memory. In general, they are very accurate for long, yet finite block lengths. It is worth mentioning that the subject of error bounds has also appeared in many more recent studies, with the advent of new approaches such as dispersion, the uncertainty-focusing bound, and the saddlepoint approximation [3], [4], [5], [6], [7]. This renewed interest in the performance of coded transmissions points to the timeliness of the topic under investigation.

A distinguishing feature of the approach that we wish to develop is the focus on indecomposable channels with memory and state-dependent operation. In many asymptotic frameworks, channel parameters are kept constant while the length of the codeword increases to infinity. Although this point of view leads to mathematically appealing characterizations, the resulting bounds on error probabilities do not depend on the initial or final states of the underlying communication channel. This situation can be explained through the fact that, no matter how slow the mixing time of the physical channel is, the duration of a

codeword eventually far exceeds this quantity. Still, in many practical scenarios, the service requirements imposed on a communication link forces the use of relatively short codewords, with no obvious time-scale separation between the duration of a codeword and the mixing time of the channel. In practice, the tradeoff between performance and delay encourages system designers to choose block lengths of the same order as the channel mixing time.

The mismatch between existing techniques and commonly deployed systems, together with the growing popularity of real-time applications on wireless networks, demands a novel approach where the impact of boundary conditions are preserved throughout the analysis. A suitable methodology should be able to capture both the effects of channel memory as well as the impact of the channel state at the onset of a codeword. In this article, we are interested in regimes where the block length is of the same order or smaller than the coherence time of the channel. Formally, we wish to study the scenario where the mixing time of the underlying finite-state channel is similar to the time necessary to transmit a codeword. This leads to two important phenomena. First, the state of the channel at the onset of a transmission has a significant impact on the empirical distribution of the states within a codeword transmission cycle. Second, channel dependencies may extend beyond the boundaries of individual codewords. This is in stark contrast with rapidly mixing channels where initial channel conditions have no effects on the probability of decoding failures. It also differs from block-fading models where the evolution of the channel is independent from block to block. Our proposed framework is rich enough to account for scenarios where decoding events are strongly correlated over time. Dependencies from codeword to codeword give rise to correlation in decoding failure events and can therefore greatly affect perceived service quality.

It is important to note that the standard asymptotic regimes naturally result in concentration of channel quality over a codeword transmission time. This enables the application of the law of large numbers to compute the mutual information, the convergence rate, etc. However, in non-asymptotic analysis of channels with memory, variations of the channel quality over a codeword transmission time need careful consideration. At the same time, high correlations make the channel behavior more predictable. This in turn, paves the way to implement adaptive opportunistic scheduling schemes. Our proposed methodology focuses on the rare-transition regime and provides results that are asymptotic in the number of bits but not in the elapsed time.

It is worth pointing out that the rare-transition regime, also known as the slow-mixing regime, has been studied in other contexts such as channel estimation, asymptotic filtering, and entropy rate analysis of hidden Markov processes [8], [9], [10], [11]. In this article, we examine probabilities of decoding failure, their distributions and temporal properties within the context of rare-transition regime. The purpose of

deriving upper bounds on the probabilities of decoding failure for rare transitions is to capture overall performance for systems that transmit data using block lengths on the order of the coherence time of their respective channels. More specifically, this article focuses on Gallager-type exponential bounds applied to probabilities of decoding failure in rare-transition regimes. By construction, these bounds depend explicitly on the initial and terminating channel states at the codeword boundaries. The analysis is conducted for the scenario where channel state information is available at the receiver. The ensuing results are subsequently compared to the probabilities of decoding failure obtained for a Gilbert-Elliott channel under a minimum distance decoder and a maximum-likelihood decision rule [1], [12], [13].

The potential implications of this novel framework are then discussed in terms of queueing theory. We exploit the results of the error-probability analysis in the rare-transition regime to evaluate the queueing performance of correlated channels. Particularly, we employ the derived upper bounds on the probability of decoding failure to bound the queueing performance of the system. In that regard, we show how stochastic dominance enables a tractable analysis of overall performance. The results of this evaluation are then compared with a performance characterization based on the exact probability of decoding failure for a Gilbert-Elliott channel under maximum-likelihood decoding.

II. MODELING AND EXPONENTIAL BOUNDS

We consider indecomposable finite-state channels where state transitions are independent of the input symbols. Such channels are often classified as fading models, and they have been used extensively in the information theory literature. We employ X_n and Y_n , respectively, to denote the input and output symbols at time n . The channel state that determines the channel law at time n is represented by S_{n-1} . We typically reserve capital letters for random variables, whereas lower case letters identify outcomes and values. Boldface letters are used to denote length- N sequences of random variables or outcomes. For groups of random variables, we use the common expression $P_{\cdot|\cdot}(\cdot|\cdot)$ to denote the conditional joint probability mass function, and $P_{e,\cdot|\cdot}(\cdot|\cdot)$ to denote the conditional joint probability of decoding error.

In general, the conditional probability distribution governing a finite-state channel can be written as

$$\begin{aligned} P_{Y_n, S_n | X_n, S_{n-1}}(y_n, s_n | x_n, s_{n-1}) \\ = \Pr(Y_n = y_n, S_n = s_n | X_n = x_n, S_{n-1} = s_{n-1}). \end{aligned}$$

When state transitions are independent of input symbols, this conditional distribution allows the factorization

$$\begin{aligned} P_{Y_n, S_n | X_n, S_{n-1}}(y_n, s_n | x_n, s_{n-1}) \\ = P_{S_n | S_{n-1}}(s_n | s_{n-1}) P_{Y_n | X_n, S_{n-1}}(y_n | x_n, s_{n-1}). \end{aligned} \tag{1}$$




Fig. 1. The Gilbert-Elliott model is the simplest, non-trivial instantiation of a finite-state channel with memory. State evolution over time forms a Markov chain and the input-output relationship of this binary channel is governed by a state-dependent crossover probability.

Throughout, we assume that the channel statistics are homogeneous over time and the sequence $\{S_n\}$ forms a Markov chain. When dealing with finite-state channels, it is customary to use integers to denote the possible input and output symbols; in our exposition, we adhere to this convention.

The famed Gilbert-Elliott channel is the proverbial example of a channel that possesses the structure described above [12], [13]. This quintessential model is governed by a two-state Markov chain, and it is illustrated in Fig. 1. The transition probability matrix for the Gilbert-Elliott channel can be expressed as

$$\mathbf{P} = \begin{bmatrix} 1 - \alpha & \alpha \\ \beta & 1 - \beta \end{bmatrix}, \quad (2)$$

where $[\mathbf{P}]_{ij} = \Pr(S_n = j | S_{n-1} = i)$. The state-dependent input-output relationship induced by channel state $s \in \{1, 2\}$ is governed by the crossover probability ε_s , where

$$\Pr(x_n = y_n | S_{n-1} = s) = 1 - \varepsilon_s$$

$$\Pr(x_n \neq y_n | S_{n-1} = s) = \varepsilon_s.$$

For convenience, we order states such that $\varepsilon_1 < \varepsilon_2$.

A coding strategy that has proven exceptionally fruitful in information theory is the use of random codes. Continuing this tradition, we adopt a random coding scheme that employs a code ensemble \mathcal{C} with $M = e^{NR}$ elements. Familiarity with this topic may help because random coding arguments tend to be notationally heavy [2]. The variable N denotes the block length of the code and R is the code rate in nats per code bit. Every element in \mathcal{C} corresponds to a sequence of admissible channel inputs, $\mathbf{x} = (x_1, x_2, \dots, x_N)$. Moreover, codewords are indexed by $k \in \{1, \dots, M\}$. The input sequence associated with index k , which we denote by $\mathbf{X}(k)$, is determined through the following procedure.

Suppose that $Q(\cdot)$ is a distribution on the set of admissible input symbols. Let $\mathbf{Q}_N(\mathbf{x}) = \prod_{n=1}^N Q(x_n)$ be the product measure induced by Q . Codeword $\mathbf{X}(k)$ is selected at random according to distribution \mathbf{Q}_N , i.e.,

$$\Pr(\mathbf{X}(k) = \mathbf{x}) = \mathbf{Q}_N(\mathbf{x}) = \prod_{n=1}^N Q(x_n).$$

We emphasize that every codeword is selected independently from other elements in \mathcal{C} . Once a code ensemble has been generated, a message is sent to the destination by first selecting one of the codewords, and then sequentially transmitting its entries over the communication channel. For the sake of clarity, we summarize our assumptions below; they apply from this point forward, unless otherwise stated.

Assumption 1: Communication takes place over a finite-state channel that admits the conditional decomposition of (1). Information is transmitted using the random coding strategy outlined above. On the receiver side, a maximum-likelihood decision rule is used to decode the received sequence. Furthermore, the state of the channel is causally known at the receiver.

At this point, it is worth restating our objective. We want to upper bound the probability that a codeword is decoded erroneously at the receiver. Concurrently, we wish to develop a rare-transition regime that remains true to the fact that the channel state at the onset of the codeword transmission process affects the evolution of the system. Ultimately, this can be achieved in an asymptotic framework by slowing down the transition profile of the underlying channel as the block length of the code grows unbounded. One of the repercussions of this setting is that we have to modify some of the results on error exponents presented by Gallager [2]. In particular, we need the ability to restrict a channel sequence \mathbf{S} to specific events and decompose the error probability accordingly.

Our first formal result is a straightforward extension to Theorem 5.6.1 in [2, pp. 135], which is itself quite general. Since we are interested in finite-state channels with memory in a slow transition regime, we require the ability to track channel realizations explicitly. From an abstract perspective, conditioning on a specific fading realization is equivalent to altering the statistical profile of the underlying channel.

Proposition 1: Suppose that the realization of the channel over the duration of a codeword is given by \mathbf{s} . Then, for any $\rho \in [0, 1]$, the probability of decoding failure at the destination, conditioned on state sequence $\mathbf{S} = \mathbf{s}$, is upper bounded by

$$P_{e|\mathbf{S}}(\mathbf{s}) \leq e^{-N(E_{0,N}(\rho, \mathbf{Q}_N, \mathbf{s}) - \rho R)}$$

where the exponent $E_{0,N}(\rho, \mathbf{Q}_N, \mathbf{s})$ is equal to

$$-\frac{1}{N} \ln \prod_{n=1}^N \sum_{y_n} \left[\sum_{x_n} Q(x_n) P_{Y_n|X_n, S_{n-1}}(y_n | x_n, s_{n-1})^{\frac{1}{1+\rho}} \right]^{1+\rho}.$$

Proof: First, we emphasize that the condition $\mathbf{S} = \mathbf{s}$ simply alters the probability measure governing the input-output relationship of the channel. Applying Theorem 5.6.1 in [2] with $M = e^{NR}$, we immediately get

$$P_{e|\mathbf{S}}(\mathbf{s}) \leq e^{\rho NR} \sum_{\mathbf{y}} \left[\sum_{\mathbf{x}} \mathbf{Q}_N(\mathbf{x}) P_{\mathbf{Y}|\mathbf{X},\mathbf{S}}(\mathbf{y}|\mathbf{x}, \mathbf{s})^{\frac{1}{1+\rho}} \right]^{1+\rho},$$

where $P_{\mathbf{Y}|\mathbf{X},\mathbf{S}}(\mathbf{y}|\mathbf{x}, \mathbf{s})$ represents the conditional distribution of receiving \mathbf{y} given $\mathbf{X} = \mathbf{x}$ and $\mathbf{S} = \mathbf{s}$. Moving forward, the crux of the argument is based on interchanging the order of exhaustive products and sums. Under the channel decomposition introduced in (1), the double summation that appears in this upper bound becomes

$$\begin{aligned} & \sum_{\mathbf{y}} \left[\sum_{\mathbf{x}} \prod_{n=1}^N Q(x_n) P_{Y_n|X_n, S_{n-1}}(y_n|x_n, s_{n-1})^{\frac{1}{1+\rho}} \right]^{1+\rho} \\ &= \prod_{n=1}^N \sum_{y_n} \left[\sum_{x_n} Q(x_n) P_{Y_n|X_n, S_{n-1}}(y_n|x_n, s_{n-1})^{\frac{1}{1+\rho}} \right]^{1+\rho}. \end{aligned}$$

Collecting these various results and using equivalent notation, we obtain the desired proposition. \blacksquare

A key insight revealed through the proof of Proposition 1 is that $E_{0,N}(\rho, \mathbf{Q}_N, \mathbf{s})$ only depends on \mathbf{s} through its empirical distribution, designated $\mathcal{T}(\mathbf{s})$. We state this fact as a corollary because it will become very useful shortly.

Corollary 1: Let T be the empirical state distribution of a sequence of N consecutive channel realizations. If \mathbf{s} and \mathbf{s}' are two sequences such that $\mathcal{T}(\mathbf{s}) = \mathcal{T}(\mathbf{s}') = T$, then with some abuse of notation we can write

$$E_{0,N}(\rho, \mathbf{Q}_N, \mathbf{s}) = E_{0,N}(\rho, \mathbf{Q}_N, \mathbf{s}') \triangleq E_{0,N}(\rho, \mathbf{Q}_N, T).$$

Furthermore, the probability of decoding failure at the destination, conditioned on $\mathbf{S} = \mathbf{s}$, is bounded by

$$P_{e|\mathbf{S}}(\mathbf{s}) \leq e^{-N(E_{0,N}(\rho, \mathbf{Q}_N, T) - \rho R)},$$

for any sequence \mathbf{s} with empirical distribution $\mathcal{T}(\mathbf{s}) = T$ and $\rho \in [0, 1]$.

For the problem at hand, we are especially interested in probabilities of the form $P_{e, S_N|S_0}(s_N|s_0)$. In some sense, each of these represents the probability of a decoding failure while keeping track of boundary states. In view of Corollary 1, it is natural to upper bound this quantity by partitioning the set of possible sequences according to their empirical distributions. This is accomplished below. In stating our results, we use \mathcal{T} to denote the collection of all admissible empirical channel distributions over sequences of length N .

Proposition 2: Suppose that a codeword is transmitted over a finite-state channel. The joint probability that decoding fails at the destination and $S_N = s_N$, conditioned on initial state $S_0 = s_0$, is upper bounded as follows

$$\begin{aligned} P_{e,S_N|S_0}(s_N|s_0) &\leq \sum_{T \in \mathcal{T}} P_{\mathcal{T}(\mathbf{S}), S_N|S_0}(T, s_N|s_0) \min_{\rho \in [0,1]} e^{-N(E_{0,N}(\rho, \mathbf{Q}_N, T) - \rho R)} \\ &\leq \min_{\rho \in [0,1]} \sum_{T \in \mathcal{T}} P_{\mathcal{T}(\mathbf{S}), S_N|S_0}(T, s_N|s_0) e^{-N(E_{0,N}(\rho, \mathbf{Q}_N, T) - \rho R)} \end{aligned} \quad (3)$$

where $P_{\mathcal{T}(\mathbf{S}), S_N|S_0}(T, s_N|s_0)$ represents the probability that $\mathcal{T}(\mathbf{S}) = T$ and $S_N = s_N$, given initial state $S_0 = s_0$.

Proof: The progression of this demonstration parallels an argument found in Section 5.9 of [2]. By partitioning the set of length- N sequences according to their empirical distributions, we can write

$$\begin{aligned} P_{e,S_N|S_0}(s_N|s_0) &= \sum_{\mathbf{s}} \Pr(\mathbf{S} = \mathbf{s}, S_N = s_N | S_0 = s_0) P_{e|\mathbf{S}}(\mathbf{s}) \\ &= \sum_{T \in \mathcal{T}} \sum_{\mathbf{s}: \mathcal{T}(\mathbf{s})=T} \Pr(\mathbf{S} = \mathbf{s}, S_N = s_N | S_0 = s_0) P_{e|\mathbf{S}}(\mathbf{s}) \\ &\leq \sum_{T \in \mathcal{T}} \sum_{\mathbf{s}: \mathcal{T}(\mathbf{s})=T} \Pr(\mathbf{S} = \mathbf{s}, S_N = s_N | S_0 = s_0) \min_{\rho \in [0,1]} e^{-N(E_{0,N}(\rho, \mathbf{Q}_N, \mathbf{s}) - \rho R)}. \end{aligned}$$

The inequality in this expression comes from a direct application of Proposition 1. Following our previous observation that $E_{0,N}(\rho, \mathbf{Q}_N, \mathbf{s})$ only depends on \mathbf{s} through its empirical distribution $\mathcal{T}(\mathbf{s})$, we can rewrite this upper bound as

$$P_{e,S_N|S_0}(s_N|s_0) \leq \sum_{T \in \mathcal{T}} \Pr(\mathcal{T}(\mathbf{S}) = T, S_N = s_N | S_0 = s_0) \min_{\rho \in [0,1]} e^{-N(E_{0,N}(\rho, \mathbf{Q}_N, T) - \rho R)}.$$

The first bound in Proposition 2 is equivalent to this inequality, yet it is expressed using a more concise notation. The second bound holds because the sum of non-negative minimums is upper bounded by the minimum of the summands. ■

In words, this result is obtained by first grouping channel state sequences according to types, applying an exponential upper bound on the probability of decoding failure to each group, and then taking an expectation over possible empirical distributions. From a large deviations perspective, this decomposition into summands is pertinent because it can be employed to identify the dominating behavior of the system as block length becomes increasingly large. Interestingly, the upper bound on the error probability only depends on the initial and final states of the channel through the empirical distributions.

Next, we consider a Gilbert-Elliott type channel, where the cardinality of the channel state space is \mathcal{S} .

Let sequence \mathbf{s} be fixed and recall that $s_n \in \{1, 2, \dots, S\}$. Then, by Proposition 1, we get

$$\begin{aligned} E_{0,N}(\rho, \mathbf{Q}_N, \mathbf{s}) &= -\frac{1}{N} \ln \prod_{n=0}^{N-1} \frac{1}{2^\rho} \left(\varepsilon_{s_n}^{\frac{1}{1+\rho}} + (1 - \varepsilon_{s_n})^{\frac{1}{1+\rho}} \right)^{1+\rho} \\ &= -\frac{1}{N} \sum_{i=1}^S n_i \ln \frac{1}{2^\rho} \left(\varepsilon_i^{\frac{1}{1+\rho}} + (1 - \varepsilon_i)^{\frac{1}{1+\rho}} \right)^{1+\rho} = \sum_{i=1}^S \frac{n_i}{N} b_i(\rho), \end{aligned} \quad (4)$$

where n_i is the number of visits to channel state i in sequence \mathbf{s} , n_i/N is the fraction of time spent in state i , and

$$b_i(\rho) = -\ln \frac{1}{2^\rho} \left(\varepsilon_i^{\frac{1}{1+\rho}} + (1 - \varepsilon_i)^{\frac{1}{1+\rho}} \right)^{1+\rho}. \quad (5)$$

Without loss of generality, we assume the error probabilities in different states are ordered such that

$$\epsilon_1 < \epsilon_2 < \dots < \epsilon_S \leq \frac{1}{2},$$

which implies

$$b_1(\rho) > b_2(\rho) > \dots > b_S(\rho). \quad (6)$$

A key observation from (4) is that the upper bound in (3) can be rewritten as an expectation with respect to the distribution of a weighted sum of the state occupation times. This observation is of significant importance, as it reduces the computational complexity of the bound in (3) by a great amount. The following remark precisely characterizes the statement.

Remark 1: Suppose the random variable N_i denotes the number of visits to channel state i over the duration of a codeword, and let

$$W(\rho) = \frac{1}{N} \sum_{i=1}^S N_i b_i(\rho), \quad (7)$$

designate the weighted sum of the normalized occupancy times. Then, the second upper bound on the error probability in (3) can be rewritten as

$$\min_{\rho \in [0,1]} \sum_w P_{W(\rho), S_N | S_0}(w, s_N | s_0) e^{-N(w - \rho R)}, \quad (8)$$

where $P_{W(\rho), S_N | S_0}(w, s_N | s_0)$ represents the probability that $W(\rho) = w$ and $S_N = s_N$, given initial state $S_0 = s_0$. We emphasize that the weights $b_i(\rho)$ are functions of the optimizing parameter ρ and are deterministic.

We note that the exponential term in (3) is averaged with respect to the joint distribution of channel state occupations $P_{\mathcal{T}(\mathbf{S}), S_N | S_0}(T, s_N | s_0)$. However, conditioned on $W(\rho)$, the bound is statistically independent of $\mathcal{T}(\mathbf{S})$. In fact, when the channel state information is available at the receiver, $W(\rho)$ is a sufficient statistic to compute the upper bound on probability of decoding error. The joint distribution of the channel

states provides more information than what is needed to derive the bound. This key characterization significantly improves the computational efficiency of the bounding technique, especially when the number of channel states increases.

For illustrative purposes, we derive the upper bounds on $P_{e,S_N|S_0}(s_N|s_0)$ for the two-state Gilbert-Elliott channel. Exploiting the Markov structure of this channel, we get

$$\begin{aligned} P_{e,S_N|S_0}(s_N|s_0) &\leq \min_{\rho \in [0,1]} \sum_{\mathbf{s}} e^{-N(E_{0,N}(\rho, \mathbf{Q}_N, \mathbf{s}) - \rho R)} \Pr(\mathbf{S} = \mathbf{s}, S_N = s_N | S_0 = s_0) \\ &= \min_{\rho \in [0,1]} \left(\mathbf{e}_{s_0} \begin{bmatrix} a(1,1) & a(1,2) \\ a(2,1) & a(2,2) \end{bmatrix}^N \mathbf{e}_{s_N}^T \right) e^{\rho NR}. \end{aligned} \quad (9)$$

In this equation, \mathbf{e}_i represents the unit vector of length two with a one in the i th position. Matrix entries are defined by $a(i,j) = [\mathbf{P}]_{ij} e^{b_i}$, where the transition probability matrix \mathbf{P} is given in (2). We emphasize that this inequality holds for any $\rho \in [0,1]$ and, hence, the bound can be tightened by minimizing over ρ .

Remark 2: The bound in (9) is very similar to Gallager's exponential bound for finite-state channels [2, Thm. 5.9.3, pp. 185] when the receiver has perfect state information. The main difference is that Gallager considers the ergodic regime and his equation simplifies to the logarithm of the largest eigenvalue of the matrix. We omit this simplification because we are mainly interested in non-asymptotic regimes.

III. THE RARE-TRANSITION REGIME

In a traditional setting where \mathbf{P} is kept constant, the upper bound given in (3) can be refined using the Perron-Frobenius theorem [2, pp. 184–185]. The more intriguing scenario for our purpose is the rare-transition regime where individual transition probabilities vary with N . We introduce such a dependency through the sampling of a continuous-time Markov chain (CTMC) $\mathcal{X}(\cdot)$, defined by the infinitesimal generator matrix \mathbf{Q} .

Following conventional notation, we use Ω to designate the sample space, and we represent a generic outcome by ω . Whenever needed, we use superscript ω to refer to a particular realization. For instance, $\mathcal{X}(\cdot)$ denotes the CTMC whereas $\mathcal{X}^\omega(\cdot)$ symbolizes the sample path associated with realization ω . In particular, $\mathcal{X}^\omega(\cdot)$ defines a mapping from the time interval $[0, \infty)$ to the state space $\{1, 2, \dots, \mathcal{S}\}$. The need for this notation will become manifest shortly.

Suppose $\mathcal{X}(\cdot)$ is sampled at every $1/N$ unit of time. Then, we can construct a continuous-time version of the sampled chain as follows,

$$\mathcal{X}_N(t) = \mathcal{X}\left(\frac{\lfloor Nt \rfloor}{N}\right). \quad (10)$$

Let \mathbf{P}_N represent the transition probability matrix of the sampled Markov chain given by $\mathcal{X}_N(n/N)$, $n \in \mathbb{N}$. Matrix \mathbf{P}_N is governed by \mathbf{Q} through the equation [14, Thms. 2.1.1 & 2.1.2]

$$\mathbf{P}_N = \exp\left(\frac{\mathbf{Q}}{N}\right). \quad (11)$$

We note that \mathbf{P}_N is also the transition probability matrix of the Markov chain $\mathcal{X}(t)$ for a time interval of length $1/N$,

$$[\mathbf{P}_N]_{ij} = \Pr\left(\mathcal{X}\left(\frac{1}{N}\right) = j \mid \mathcal{X}(0) = i\right).$$

In a similar fashion, we can make a distinction between the sampled chain \mathcal{X}_N and the realization \mathcal{X}_N^ω associated with outcome ω .

As an example, consider a two-state Markov process. In this case, the infinitesimal generator matrix can be written as

$$\mathbf{Q} = \begin{bmatrix} -\mu & \mu \\ \xi & -\xi \end{bmatrix} \quad \mu, \xi > 0, \quad (12)$$

and, consequently, we can express the transition probability matrix of the sampled process as

$$\mathbf{P}_N = \frac{1}{\mu + \xi} \begin{bmatrix} \xi + \mu e^{-\frac{\xi+\mu}{N}} & \mu \left(1 - e^{-\frac{\xi+\mu}{N}}\right) \\ \xi \left(1 - e^{-\frac{\xi+\mu}{N}}\right) & \mu + \xi e^{-\frac{\xi+\mu}{N}} \end{bmatrix}. \quad (13)$$

As seen above, jumps in the discrete chain become less likely when N increases. This should be expected because a refined sampling of the CTMC does not alter the character of the underlying process. Furthermore, the roles of boundary states are preserved, a property which is key for our subsequent analysis. At this point, it is instructive to note that the inequalities presented in Section II apply in the context of rare transitions as well, albeit using \mathbf{P}_N rather than a fixed \mathbf{P} .

An important benefit of the rare-transition regime is the existence of approximate error bounds that can be computed efficiently. These approximate bounds can be obtained using the following steps. First, we show that the distributions of the occupation times for the sampled Markov chains converge to the distribution of the channel state occupation times of the original CTMC, as the sampling interval $1/N$ decreases to zero. Second, we employ standard results pertaining to the convergence of empirical measures to get approximate upper bounds on the probabilities of decoding failure at the destination, using the continuous measures. We then leverage a numerical procedure to compute the distributions of weighted sums of channel state occupations for the CTMC [15], [16]. Collecting these results, we arrive at the desired characterization of channels with memory.

Recall the CTMC $\mathcal{X}(t)$ and its sampled variants $\mathcal{X}_N(t)$, as introduced in (10). For every channel state i , we define the occupation times pathwise through the integrals below,

$$\eta_i^\omega = \int_0^1 \mathbf{1}_{\{\mathcal{X}^\omega(t)=i\}} dt$$

$$\eta_{N,i}^\omega = \int_0^1 \mathbf{1}_{\{\mathcal{X}_N^\omega(t)=i\}} dt = \frac{1}{N} \sum_{k=1}^N \mathbf{1}_{\{\mathcal{X}_N^\omega(k/N)=i\}},$$

where $\mathbf{1}_{\{\cdot\}}$ denotes the standard indicator function. Having specified the values of the occupation times for every possible outcome ω , these equations unambiguously define random variables η_i and $\eta_{N,i}$.

Proposition 3: The sequence of random vectors given by

$$\bar{\eta}_N = (\eta_{N,1}, \dots, \eta_{N,S}), \quad N = 1, 2, \dots$$

converges almost surely to random vector $\bar{\eta} = (\eta_1, \dots, \eta_S)$ as N approaches infinity.

Proof: The CTMC $\mathcal{X}(t)$ is time-homogeneous and its state space has finite cardinality. It is therefore non-explosive [14, Sec. 2.7]. This implies that the number of transitions in the interval $[0, 1]$ is finite almost surely; we can then write $\Pr(\Omega') = 1$, where

$$\Omega' = \{\omega \in \Omega \mid \mathcal{X}^\omega(t) \text{ has finitely many jumps in } [0, 1]\}.$$

For any $\omega \in \Omega'$, the function \mathcal{X}^ω is bounded and continuous almost everywhere on $[0, 1]$. It therefore fulfills Lebesgue's criterion for Riemann integrability [17, pp. 323] and, as such,

$$\begin{aligned} \eta_i^\omega &= \int_0^1 \mathbf{1}_{\{\mathcal{X}^\omega(t)=i\}} dt \\ &= \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{k=1}^N \mathbf{1}_{\{\mathcal{X}_N^\omega(k/N)=i\}} = \lim_{N \rightarrow \infty} \eta_{N,i}^\omega. \end{aligned}$$

Since the number of channel states is finite, this result readily extends to vectors,

$$\lim_{N \rightarrow \infty} d_1(\bar{\eta}_N^\omega, \bar{\eta}^\omega) = 0 \quad \forall \omega \in \Omega',$$

where $d_1(\cdot, \cdot)$ is the ℓ_1 distance on \mathbb{R}^S . Equivalently, we can write

$$\Pr\left(\omega \in \Omega \mid \lim_{N \rightarrow \infty} d_1(\bar{\eta}_N^\omega, \bar{\eta}^\omega) = 0\right) = 1.$$

That is, $\bar{\eta}_N$ converges to $\bar{\eta}$ almost surely, as desired. ■

Almost sure convergence implies convergence in probability and in distribution [18, Sec. 8.5]. Thus, from Proposition 3, we gather that $\bar{\eta}_N$ converges to $\bar{\eta}$ in distribution, which is sufficient for our later purpose. This further implies that the occupation times of a sequence of independent discrete-time Markov

chains, each generated according to \mathbf{P}_N , converge in distribution to $\bar{\eta}$. That is, the sampling of the CTMC described above is a powerful mathematical tool to get the result we seek.

In the discussion above, we are intentionally vague about the probability laws on Ω . The careful reader will note that our arguments apply to the equilibrium distribution of the Markov chain as well as the conditional measure where the Markov process starts in state $i \in \{1, 2, \dots, S\}$, at time zero. Moreover, by extension, these findings apply to probabilities where the final channel state is taken into account. To distinguish between these different scenarios, we introduce a shorthand notation for these joint probabilities,

$$F(r_1, r_2, \dots, r_{S-1}) = \Pr(\eta_1 \leq r_1, \dots, \eta_{S-1} \leq r_{S-1})$$

$$F_{ij}(r_1, r_2, \dots, r_{S-1}) = \Pr(\eta_1 \leq r_1, \dots, \eta_{S-1} \leq r_{S-1}, S_f = j, S_i = i)$$

$$F_{j|i}(r_1, r_2, \dots, r_{S-1}) = \Pr(\eta_1 \leq r_1, \dots, \eta_{S-1} \leq r_{S-1}, S_f = j | S_i = i).$$

In our labeling, S_i identifies the initial state of the channel and S_f specifies its final value. We can define F_N , $F_{N,ij}$ and $F_{N,j|i}$ in an analogous manner. It is immediate from Proposition 3 that $dF_N \Rightarrow dF$, $dF_{N,ij} \Rightarrow dF_{ij}$, and $dF_{N,j|i} \Rightarrow dF_{j|i}$ as N grows to infinity, where the symbol \Rightarrow denotes convergence in distribution. As a direct consequence of Proposition 3, we can apply these results to affine combinations of η_1, \dots, η_S .

Corollary 2: Let ρ be fixed and recall the coefficients $b_i(\rho)$ found in (5). The sequence of random variables given by

$$W_N(\rho) = \sum_{i=1}^S \eta_{N,i} b_i(\rho)$$

converges in distribution to random variable

$$W(\rho) = \sum_{i=1}^S \eta_i b_i(\rho)$$

as N approaches infinity.

We point out that the expression for $W(\rho)$ above differs slightly from (7), as it reflects the notation developed for the current asymptotic setting. Again, this result is valid for the probability laws associated with F , F_{ij} and $F_{j|i}$. The weighted sum $W(\rho)$ is of such importance in our impending discussion that

we introduce a convenient notation for its corresponding probability laws as well,

$$G(w) = \Pr(W(\rho) \leq w)$$

$$G_{ij}(w) = \Pr(W(\rho) \leq w, S_i = i, S_f = j)$$

$$G_{j|i}(w) = \Pr(W(\rho) \leq w, S_f = j | S_i = i).$$

In addition, we write G_N , $G_{N,ij}$ and $G_{N,j|i}$ for the measures associated with $W_N(\rho)$. It may be helpful to point out that we introduce a slight abuse of notation in establishing these quantities; the dependence of these probability laws on ρ is implicit. This is intentional as the alternative makes the notation overly cumbersome and confusing. In the limiting case, the measures G , G_{ij} , and $G_{j|i}$ are continuous almost everywhere. Wherever needed, we can emphasize the dependence on ρ by writing

$$dG(w) = f_{W(\rho)}(w)dw$$

$$dG_{ij}(w) = f_{W(\rho), S_i, S_f}(w, i, j)dw$$

$$dG_{j|i}(w) = f_{W(\rho), S_f | S_i}(w, j|i)dw$$

where $f(\cdot)$ is our generic representation of a probability density, possibly with weighted Dirac delta components.

Proposition 4: Suppose that a message is transmitted over a fading channel with \mathcal{S} states, using the random coding scheme introduced in Section II. An approximate upper bound for the error probability $P_{e, S_N | S_0}(j|i)$ is given by

$$P_{e, S_N | S_0}(j|i) \lesssim \int_{b_S(\rho)}^{b_1(\rho)} \min \left\{ 1, e^{-N(w-\rho R)} \right\} f_{W(\rho), S_f | S_i}(w, j|i)dw. \quad (14)$$

The approximation in (14) reflects its potential use in selecting efficient coding schemes. The precise mathematical meaning underlying this equation is described in the proof.

Proof: Let $\rho \in (0, 1)$ be fixed. We know from Corollary 1 that, for channel type T , the error probability is bounded by

$$P_{e|S}(T) \leq e^{-N(E_{0,N}(\rho, \mathbf{Q}_N, T) - \rho R)} = e^{-N(w - \rho R)}.$$

where $w = \sum_{i=1}^S \frac{n_i}{N} b_i(\rho)$ is determined by the channel type. We can readily tighten this bound to

$$P_{e|S}(T) \leq \min \left\{ 1, e^{-N(w - \rho R)} \right\}$$

because individual probabilities cannot exceed one. It is useful to point out that the expression $w - \rho R$ is an affine, strictly increasing function of w . For the purpose of exposition, let $g_N(w)$ be defined by

$$g_N(w) = \min \left\{ 1, e^{-N(w-\rho R)} \right\} = \begin{cases} e^{-N(w-\rho R)} & w < \varrho \\ 1 & w \geq \varrho \end{cases},$$

where the threshold $\varrho = \rho R$. The sequence of functions $\{g_N(w)\}$ converges pointwise to

$$g(w) = \mathbf{1}_{[\varrho, \infty)}(w) = \begin{cases} 0 & w < \varrho \\ 1 & w \geq \varrho \end{cases},$$

which is uniformly continuous on the set $w \in [0, b_1(\rho)] \setminus \{\varrho\}$. By taking an expectation over $W(\rho)$, we get

$$\begin{aligned} P_{e, S_N | S_0}(j|i) &\leq \int_{b_S(\rho)}^{b_1(\rho)} \min \left\{ 1, e^{-N(w-\rho R)} \right\} dG_{N,j|i}(w) \\ &= \sum_w P_{W_N(\rho), S_N | S_0}(w, j|i) \min \left\{ 1, e^{-N(w-\rho R)} \right\}. \end{aligned} \quad (15)$$

While this upper bound can be computed numerically if the distribution of $W_N(\rho)$ is known, we are also interested in approximations that provide good intuition as well as computational efficiency at the expense of a little accuracy.

As a next step, we will establish that

$$\lim_{N \rightarrow \infty} \int_{b_S(\rho)}^{b_1(\rho)} g_N dG_{N,j|i} = \int_{b_S(\rho)}^{b_1(\rho)} g dG_{j|i}. \quad (16)$$

From Corollary 2, we know that the set of measures $\{G_{N,j|i}\}$ converges in distribution to $G_{j|i}$. Also, for any converging sequence $w_N \in [b_S(\rho), b_1(\rho)]$ with $\lim_{N \rightarrow \infty} w_N = w \neq \varrho$, we have $g_N(w_N) \rightarrow g(w)$. This is pertinent because the limiting measure $G_{j|i}(w)$ is continuous at $\varrho \in (b_S(\rho), b_1(\rho))$ and, consequently, the event $\{w = \varrho\}$ has probability zero. Collecting these observations, we can apply [19, Thm. 5.5] and thereby establish the validity of (16).

In view of these results, we can write

$$\int_{b_S(\rho)}^{b_1(\rho)} g_N dG_{N,j|i} \approx \int_{b_S(\rho)}^{b_1(\rho)} g_N dG_{j|i} \approx \int_{b_S(\rho)}^{b_1(\rho)} g dG_{j|i} \quad (17)$$

for large enough values of N . Since the first integral in (17) provides an upper bound on $P_{e, S_N | S_0}(j|i)$, the approximate upper bound of (14) immediately follows. That is, if the code length is large enough, then the approximation in Proposition 4 is justified. ■

Using (14) to select system parameters gives an alternate, computationally efficient way to design good systems. We emphasize that one does not necessarily need to compute a sequence of distributions for

$\{W_N(\rho)\}$ to follow this solution path. Rather, it is possible to accurately approximate the distribution of $W(\rho)$ directly using an iterative approach. We review one such method below; it leverages numerical techniques introduced in [15], [16] to compute the distributions of reward functions on CTMCs. This method applies to channels with arbitrary, yet finite numbers of states. In contrast, the standard approach associated with (15) entails computing $G_{N,j|i}$ explicitly for multiple values of N , a cumbersome task.

Proposition 5 presents a numerical method to compute the distribution of $W(\rho)$. This method is adapted from [15], [16] and, as such, it is presented without detailed proof. In practice, the infinite sum needs to be truncated according to an appropriate criterion. To present this result, we need to introduce relevant notations. Let \mathcal{A} be the transition matrix given by

$$\mathcal{A} = \mathbf{I} + \frac{\mathbf{Q}}{\sigma}$$

where \mathbf{I} is the identity matrix, $\sigma \geq \max_k |\mathbf{Q}_{kk}|$, $k \in \{1, 2, \dots, \mathcal{S}\}$, is a constant and $\{\mathbf{Q}_{kk}\}$ are the diagonal elements of \mathbf{Q} . Also, define matrix $\mathbf{G}(w)$ by

$$[\mathbf{G}(w)]_{ij} = G_{j|i}(w)$$

where $i, j \in \{1, 2, \dots, \mathcal{S}\}$.

Proposition 5: Let ρ be fixed and suppose that the channel is initially in state $S_i = i$. The probabilities of the events $\{W(\rho) \leq w, S_f = j | S_i = i\}$ as functions of w are continuous almost everywhere, and they have at most \mathcal{S} discontinuities, with possible locations $b_{\mathcal{S}}(\rho), \dots, b_1(\rho)$. Furthermore, for $w \in [b_k(\rho), b_{k-1}(\rho))$ and $2 \leq k \leq \mathcal{S}$, we have

$$\mathbf{G}(w) = \sum_{n=0}^{\infty} e^{-\sigma} \frac{\sigma^n}{n!} \sum_{l=0}^n \binom{n}{l} w_k^l (1 - w_k)^{n-l} \mathbf{C}^{(k)}(n, l),$$

where

$$w_k = \frac{w - b_k(\rho)}{b_{k-1}(\rho) - b_k(\rho)}.$$

The matrices $\{\mathbf{C}^{(k)}(n, l)\}$ are defined component-wise by

$$[\mathbf{C}^{(k)}(n, l)]_{cd} = C_{cd}^{(k)}(n, l) \quad c, d \in \{1, 2, \dots, \mathcal{S}\},$$

and the individual entries in each of these matrices are given by the following two recurrence relations.

For $k \leq c \leq \mathcal{S}$ and $0 \leq d \leq \mathcal{S}$,

$$\begin{aligned} C_{cd}^{(k)}(n, l) &= \frac{b_k(\rho) - b_c(\rho)}{b_{k-1}(\rho) - b_c(\rho)} C_{cd}^{(k)}(n, l-1) \\ &+ \frac{b_{k-1}(\rho) - b_k(\rho)}{b_{k-1}(\rho) - b_c(\rho)} \sum_{e=0}^{\mathcal{S}} [\mathcal{A}]_{ce} C_{ed}^{(k)}(n-1, l-1) \end{aligned}$$

where $1 \leq l \leq n$. For $n \geq 0$ and $k > 1$, we apply the boundary conditions $C_{cd}^{(1)}(n, 0) = 0$ and $C_{cd}^{(k)}(n, 0) = C_{cd}^{(k-1)}(n, n)$.

Similarly, for $0 \leq c \leq k-1$ and $0 \leq d \leq \mathcal{S}$,

$$C_{cd}^{(k)}(n, l) = \frac{b_c(\rho) - b_{k-1}(\rho)}{b_c(\rho) - b_k(\rho)} C_{cd}^{(k)}(n, l+1) + \frac{b_{k-1}(\rho) - b_k(\rho)}{b_k(\rho) - b_c(\rho)} \sum_{e=0}^{\mathcal{S}} [\mathcal{A}]_{ce} C_{ed}^{(k)}(n-1, l)$$

where $0 \leq l \leq n-1$. In this case, for $n \geq 0$ and $k < \mathcal{S}$, we can write the boundary conditions $C_{cd}^{(\mathcal{S})}(n, n) = [\mathcal{A}^n]_{cd}$ and $C_{cd}^{(k)}(n, n) = C_{cd}^{(k+1)}(n, 0)$.

Sketch of proof: We emphasize, again, that this proposition is adapted from a general technique found in [15], [16]. In paralleling the argument presented therein, the weights $b_1(\rho), \dots, b_{\mathcal{S}}(\rho)$ play the role of reward rates and $W(\rho)$ represents the total continuous reward over the interval $[0, 1)$. The possible discontinuities in $G_{j|i}(w)$ have to do with the non-vanishing probabilities that the chain does not visit certain states during time interval $[0, 1)$. ■

The discrete-time Markov chain whose probability transition matrix is given by \mathcal{A} is called a uniformized chain [20], [21]. This chain can be paired to a Poisson process with rate σ to form a continuous-time Markov chain. The resulting chain is stochastically equivalent to $\mathcal{X}(t)$ and, as such, it possesses the same invariant probability distribution [15]. Furthermore, the matrix \mathbf{P}_N can be written as

$$\mathbf{P}_N = e^{-\frac{\sigma}{N}} \exp\left(\frac{\sigma \mathcal{A}}{N}\right),$$

or, alternatively,

$$[\mathbf{P}_N]_{ij} = e^{-\frac{\sigma}{N}} \sum_{k=0}^{\infty} \frac{1}{k!} \left(\frac{\sigma}{N}\right)^k [\mathcal{A}^k]_{ij}.$$

While the uniformized chain and the sampled chain are both derived from \mathbf{Q} , there remains an important difference. By construction, the uniformized chain never misses a transition in its corresponding continuous-time Markov chain; whereas the sampled chain with its periodic structure can overlook jumps associated with fast transitions. This makes the uniformized chain a more suitable object in describing Proposition 5.

As a special case of Proposition 5, we turn to the situation where $\mathcal{S} = 2$. Not surprisingly, occupation times for this simple scenario have been studied in the past, and explicit expressions for their distributions exist [22], [23], [24]. Still, the distributions provided therein only account for an initial state, and they do not specify a final state. We must therefore modify these results slightly to match the needs of our current framework.

Lemma 1: Consider a continuous-time Markov chain whose generator matrix is given by (12). The joint distributions governing occupation times and the final state, conditioned on the initial state, can be written as follows

$$\begin{aligned} f_{\eta_i, S_f | S_i}(r, 1|1) &= e^{-\mu r - \xi(1-r)} \left(\delta(1-r) + \sqrt{\frac{\mu \xi r}{1-r}} I_1 \left(2\sqrt{\mu \xi r(1-r)} \right) \right) \\ f_{\eta_i, S_f | S_i}(r, 2|1) &= \mu e^{-\mu r - \xi(1-r)} I_0 \left(2\sqrt{\mu \xi r(1-r)} \right) \\ f_{\eta_i, S_f | S_i}(r, 1|2) &= \xi e^{-\mu r - \xi(1-r)} I_0 \left(2\sqrt{\mu \xi r(1-r)} \right) \\ f_{\eta_i, S_f | S_i}(r, 2|2) &= e^{-\mu r - \xi(1-r)} \left(\delta(r) + \sqrt{\frac{\mu \xi (1-r)}{r}} I_1 \left(2\sqrt{\mu \xi r(1-r)} \right) \right) \end{aligned}$$

where $I_0(\cdot)$ and $I_1(\cdot)$ represent modified Bessel functions of the first kind defined in [23, Lem. 2, pp. 386].

Proof: See the appendix. ■

The corresponding expressions for the sampled chain are presented in the following Lemma.

Lemma 2: Consider a two-state channel whose transition probability matrix is given by (2). Assume that the number of visits to each state is recorded for a period spanning N consecutive channel realizations. The joint distributions governing the channel type and its final state, conditioned on the initial state, can be written in terms of the Gaussian hypergeometric function ${}_2F_1(\cdot, \cdot; \cdot; \cdot)$ defined in [23, Lem. 1, pp. 383]. In particular, for $m = 1, 2, \dots, N-1$, we get

$$\begin{aligned} P_{N_1, S_N | S_0}(m, 1|1) &= (1-\alpha)^m (1-\beta)^{N-m} \left({}_2F_1(-N+m, -m; 1; \lambda) - {}_2F_1(-N+m+1, -m; 1; \lambda) \right) \\ P_{N_1, S_N | S_0}(m, 2|1) &= \frac{(1-\alpha)^{m-1} (1-\beta)^{N-m+1} \alpha}{(1-\beta)} {}_2F_1(-N+m, -m+1; 1; \lambda) \\ P_{N_1, S_N | S_0}(m, 1|2) &= \frac{(1-\alpha)^{m+1} (1-\beta)^{N-m-1} \beta}{(1-\alpha)} {}_2F_1(-N+m+1, -m; 1; \lambda) \\ P_{N_1, S_N | S_0}(m, 2|2) &= (1-\alpha)^m (1-\beta)^{N-m} \left({}_2F_1(-N+m, -m; 1; \lambda) - {}_2F_1(-N+m, -m+1; 1; \lambda) \right) \end{aligned}$$

where N_1 is the number of visits to the first state, and $\lambda = \frac{\alpha\beta}{(1-\alpha)(1-\beta)}$. Special consideration must be given to extremal cases. In particular, we have

$$\begin{aligned} P_{N_1, S_N | S_0}(0, \cdot|1) &= P_{N_1, S_N | S_0}(N, \cdot|2) = 0 \\ P_{N_1, S_N | S_0}(0, 2|2) &= (1-\beta)^N \\ P_{N_1, S_N | S_0}(N, 1|1) &= (1-\alpha)^N. \end{aligned}$$

Proof: See the appendix. ■

We can relate these two results through Proposition 3 and the definition of \mathbf{P}_N in (13). Let α and β be the constants defined in Lemma 2, and consider the assignment

$$\alpha = \frac{\mu}{\mu + \xi} \left(1 - e^{-\frac{\xi + \mu}{N}}\right)$$

$$\beta = \frac{\xi}{\mu + \xi} \left(1 - e^{-\frac{\xi + \mu}{N}}\right).$$

Then, the discrete distributions specified above converge to the occupation times described in Lemma 1, under the asymptotic scaling $\frac{N_1}{N} \rightarrow \eta$, as N grow unbounded. As a side observation, we emphasize that two of the limiting measures in Lemma 1 are not absolutely continuous with respect to the Lebesgue measure. Still, these distributions are well-defined positive measures [18, Chap. 8]. The discontinuities can be explained by the fact that, in the rare-transition limit, there is a positive probability of staying in the initial channel state for the whole block.

Using the limiting distributions in Lemma 1, one can compute the approximate upper bound found in (14) for the two-state case. We stress that, for this simple channel, $W(\rho) = (b_1(\rho) - b_2(\rho))\eta + b_2(\rho)$ is an affine function of η . Hence, the distribution of $W(\rho)$ can be derived in terms of η . For this simple case, it is straightforward to compute the approximate bound using the empirical distribution of the state occupancies,

$$P_{e, S_N | S_0}(j|i) \lesssim \int_0^1 \min \left\{ 1, e^{-N((b_1(\rho) - b_2(\rho))r + b_2(\rho) - \rho R)} \right\} f_{\eta, S_f | S_i}(r, j|i) dr. \quad (18)$$

Yet, as the number of states increases, dealing with the joint distribution of the state occupation times and integrating over multiple variables is an increasingly complex task. This difficulty is bypassed when we use the distribution of the weighted sum of occupation times since, in this latter case, we are dealing with a single random variable as opposed to a random vector.

In general, getting the distributions of the occupation times for the discrete chains is not needed to apply the result of Proposition 4. However, for the two-state channel, the distributions are available for both the continuous-time and the sampled chains. It is then instructive to compute the exact upper bound in (15) using the distributions given by Lemma 2, and compare it to the approximate bound presented in (18). Numerical results for this comparison are presented in Section VIII, thereby offering supporting evidence for our proposed methodology. We note that, even for the simple two-state case, computing the approximate upper bound in (18) is considerably more efficient than calculating (15). As we will see in Section VIII, the price to pay for this computational efficiency is a small loss in accuracy.

From an engineering point of view, we are interested in cases where N is dictated by the code length of a practical coding scheme. The approximate upper bound can be used to perform a quick survey of good

parameters. Then, the exact expression based on the hypergeometric function can be employed for fine tuning locally. As a final note on this topic, we emphasize that these upper bounds can be tightened by optimizing over $\rho \in [0, 1]$. This task entails repeated computations of the bounds, which partly explains our concerns with computational efficiency.

A significant benefit in dealing with the Gilbert-Elliott channel model is its tractability. The remaining of this article is devoted to the analysis of error probability and the queueing behavior of systems built around this two-state channel. The next sections are dedicated to exact derivation of probabilities of detected and undetected decoding failure. This is an intermediate step to characterize system performance, and it allows for fair evaluation of the proposed bounding technique. Furthermore, we also show how to bound the probability of undetected errors with slight modifications to the approximated upper bounds. In Sections VI and VII, we explore the queueing performance of the system, using exact expressions and upper bounds on the detected and undetected probabilities of decoding failure.

IV. EXACT PROBABILITY OF DECODING FAILURES

It is possible to compute exact probabilities of decoding failure under various decision schemes for the two-state Gilbert-Elliott channel, due to its simplicity. Consequently, in this case, we can assess how close the bounds and the true probabilities of error are from one another. The hope is that, if the Gilbert-Elliott bounds are reasonably tight, then the upper bounds for general finite-state channels will also be good. Of course, it may be impractical to compute exact probabilities of error for more elaborate channels. Even for Gilbert-Elliott type channels with more than two states, deriving and computing exact expressions for the probability of decoding failure rapidly becomes intractable. In such situations, the use of upper bounds for performance evaluation is inevitable.

In [25], the authors study data transmission over a Gilbert-Elliott channel using random coding when the state is known at the receiver. Two different decoding schemes are considered: a minimum-distance decoder and a maximum-likelihood decision rule. For the sake of completeness, we briefly review these results. When channel state information is available at the destination, the empirical distribution of the channel sequence provides enough information to determine the probability of decoding failure. Using the measure on N_1 and the corresponding conditional error probabilities, one can average over all possible types to get the probability of decoding failure,

$$P_{e,S_N|S_0}(s_N|s_0) = \sum_{T \in \mathcal{T}} P_{e|\mathcal{T}(\mathbf{S})}(T) \Pr(\mathcal{T}(\mathbf{S}) = T, S_N = s_N | S_0 = s_0). \quad (19)$$

The conditional probability of decoding failure, given type $T = (n_1, N - n_1)$, is examined further below. The probability distributions governing different channel types can be found in Lemma 2.

Consider the channel realization over the span of a codeword. Suppose \mathbf{X}_i and \mathbf{Y}_i represent the subvectors of \mathbf{X} and \mathbf{Y} corresponding to time instants when the channel is in state i . We can denote the number of errors that occur in each state using random variables E_1 and E_2 , where $E_i = d_H(\mathbf{X}_i, \mathbf{Y}_i)$ and $d_H(\cdot, \cdot)$ is the Hamming distance. The conditional error probabilities can then be written as

$$P_{e|\mathcal{T}(\mathbf{s})}(T) = \sum_{e_1=0}^{n_1} \sum_{e_2=0}^{n_2} P_{e|\mathcal{T}(\mathbf{s}), E_1, E_2}(T, e_1, e_2) P_{E_1, E_2|\mathcal{T}(\mathbf{s})}(e_1, e_2|T), \quad (20)$$

where $n_2 = N - n_1$. Given the channel type, the numbers of errors in the good and bad states are independent,

$$P_{E_1, E_2|\mathcal{T}(\mathbf{s})}(e_1, e_2|T) = P_{E_1|\mathcal{T}(\mathbf{s})}(e_1|T) P_{E_2|\mathcal{T}(\mathbf{s})}(e_2|T). \quad (21)$$

Furthermore, E_1 and E_2 have binomial distributions

$$P_{E_i|\mathcal{T}(\mathbf{s})}(e_i|T) = P_{E_i|N_i}(e_i|n_i) = \binom{n_i}{e_i} \varepsilon_i^{e_i} (1 - \varepsilon_i)^{n_i - e_i}. \quad (22)$$

This mathematical structure leads to the following theorem.

Theorem 1: When ties are treated as errors, the probability of decoding failure for a length- N uniform random code with M codewords, conditioned on the number of symbol errors in each state and the channel state type, is given by

$$P_{e|\mathcal{T}(\mathbf{s}), E_1, E_2}(T, e_1, e_2) = 1 - \left(1 - 2^{-N} \sum_{(\tilde{e}_1, \tilde{e}_2) \in \mathcal{M}(\gamma e_1 + e_2)} \binom{n_1}{\tilde{e}_1} \binom{n_2}{\tilde{e}_2} \right)^{M-1} \quad (23)$$

where $\mathcal{M}(d)$ is the set of pairs $(\tilde{e}_1, \tilde{e}_2) \in \{0, \dots, N\}^2$ that satisfy $\gamma \tilde{e}_1 + \tilde{e}_2 \leq d$. This expression holds with

$$\gamma = \frac{\ln \varepsilon_1 - \ln(1 - \varepsilon_1)}{\ln \varepsilon_2 - \ln(1 - \varepsilon_2)}$$

for the maximum-likelihood decision rule, and with $\gamma = 1$ for minimum-distance decoding.

Proof: See [25]. ■

If the channel state is causally known at the receiver, then random codes paired with a maximum-likelihood decoding rule form a permutation invariant scheme. The decoding performance is then determined by the number of symbol errors in each state within a codeword, and not by their order or locations. We point out that the methodology introduced herein can potentially be extended to other permutation invariant encoding/decoding schemes to analyze probability of decoding failure.

V. UNDETECTED ERRORS

A serious matter with pragmatic communication systems is the presence of undetected decoding failures. In the current setting, this occurs when the receiver uniquely decodes to the wrong codeword. For delay-sensitive applications, this problem is especially important because recovery procedures can lead to undue delay. To address this issue, we apply techniques that help control the probability of admitting erroneous codewords [26], [27]. This safeguard, in turn, leads to slight modifications to the performance analysis presented above. In applying these techniques, the probability of undetected failure is a system parameter that must be set during the design phase of the system.

A. The Exact Approach

In [25], the authors show that the probability of decoding failure (including detected errors, undetected errors, and ties) is given by the equations (19)–(22) and substituting

$$P_{e|\mathcal{T}(\mathbf{S}),E_1,E_2}(T, e_1, e_2) = 1 - \left(1 - 2^{-N} \sum_{(\tilde{e}_1, \tilde{e}_2) \in \mathcal{M}(\gamma e_1 + e_2 + \nu)} \binom{n_1}{\tilde{e}_1} \binom{n_2}{\tilde{e}_2} \right)^{M-1}, \quad (24)$$

where ν is a non-negative parameter that specifies the size of the safety margin for undetected errors. The joint probability of undetected error with ending state S_N , conditioned on starting in state S_0 , is $P_{\text{ue}, S_N|S_0}(s_N|s_0)$ and can be upper bounded by

$$\bar{P}_{\text{ue}, S_N|S_0}(s_N|s_0) = \sum_{T \in \mathcal{T}} \bar{P}_{\text{ue}|\mathcal{T}(\mathbf{S})}(T) \Pr(\mathcal{T}(\mathbf{S}) = T, S_N = s_N | S_0 = s_0) \quad (25)$$

where

$$\bar{P}_{\text{ue}|\mathcal{T}(\mathbf{S})}(T) = \sum_{e_1=0}^{n_1} \sum_{e_2=0}^{n_2} \bar{P}_{\text{ue}|\mathcal{T}(\mathbf{S}),E_1,E_2}(T, e_1, e_2) P_{E_1,E_2|\mathcal{T}(\mathbf{S})}(e_1, e_2|T), \quad (26)$$

and

$$\bar{P}_{\text{ue}|\mathcal{T}(\mathbf{S}),E_1,E_2}(T, e_1, e_2) = 1 - \left(1 - 2^{-N} \sum_{(\tilde{e}_1, \tilde{e}_2) \in \mathcal{M}(\gamma e_1 + e_2 - \nu)} \binom{n_1}{\tilde{e}_1} \binom{n_2}{\tilde{e}_2} \right)^{M-1}. \quad (27)$$

Since the probability of an undetected error is typically much smaller than that of a detected error, one can upper bound the probability of detected error by $P_{e, S_N|S_0}(s_N|s_0)$ with a negligible penalty.

B. Exponential Bound

With slight modifications to the derived exponential upper bound, one can get a similar bound on the probability of undetected error.

Lemma 3: The exponential upper bounds on $P_{e,S_N|S_0}$, and $\bar{P}_{ue,S_N|S_0}$ can be written as

$$\begin{aligned}\tilde{P}_{ue,S_N|S_0}(j|i) &= \min_{0 \leq v \leq \rho \leq 1} \sum_{n_1=0}^N \min \left\{ 1, e^{-N(E_{0,N}(\rho, \mathbf{Q}_N, n_1) - \rho R - v\tau)} \right\} P_{N_1, S_N|S_0}(n_1, j|i) \\ &\approx \min_{0 \leq v \leq \rho \leq 1} \int_0^1 \min \left\{ 1, e^{-N((b_1(\rho) - b_2(\rho))r + b_2(\rho) - \rho R - v\tau)} \right\} f_{\eta, S_i|S_0}(r, j|i) dr\end{aligned}\quad (28)$$

and

$$\begin{aligned}\tilde{P}_{e,S_N|S_0}(j|i) &= \min_{0 \leq v \leq \rho \leq 1} \sum_{n_1=0}^N \min \left\{ 1, e^{-N(E_{0,N}(\rho, \mathbf{Q}_N, n_1) - \rho R - v\tau + \tau)} \right\} P_{N_1, S_N|S_0}(n_1, j|i) \\ &\approx \min_{0 \leq v \leq \rho \leq 1} \int_0^1 \min \left\{ 1, e^{-N((b_1(\rho) - b_2(\rho))r + b_2(\rho) - \rho R - v\tau + \tau)} \right\} f_{\eta, S_i|S_0}(r, j|i) dr\end{aligned}\quad (29)$$

where τ controls the tradeoff between detected and undetected errors and is used to decrease the incidence of undetected errors, in a manner similar to ν for the exact case.

Proof: Following the same approach as in [26], [27], the results are achieved. ■

We emphasize that the rare-transition regime, the bounds and the approximation methodologies proposed in this article have a wide range of applications. In fact, the introduced upper bounds can be adopted in various analysis frameworks. The following two sections are dedicated to the potential implications of the proposed bounding techniques in terms of queueing theory. We exploit these results to evaluate the queueing performance of systems built around correlated channels. In particular, we show how stochastic dominance can be combined with these tools to make performance analysis tractable.

VI. QUEUEING MODEL

Consider a queueing system in which packets are generated at the source according to a Poisson process with arrival rate λ , measured in packets per channel use. The number of bits per packet forms a sequence of independent geometric random variables, each with parameter $\rho \in (0, 1)$. On arrival, a packet is divided into segments of size $RN/\ln 2$ bits, where N denotes the block length and R is the code rate in nats per code bit, as defined in Section II. The number of information bits per segment, $RN/\ln 2$, is assumed to be an integer. Since the convention for rate differs between the exponential upper bounds and the standard coding theory, we let $\tilde{R} = R/\ln 2$ be the code rate in information bits per code bit associated with the code rate R in information nats per code bit. The total number of segments associated with a packet of length L is given by $J = \left\lceil \frac{L}{\tilde{R}N} \right\rceil$. As before, a random coding scheme is used to protect the transmitted data while it is transmitted over the correlated channel.

On the receiving end, one must successfully decode all J codewords to recover the corresponding packet. Once this is achieved, this packet is discarded from the queue. We note that random variable J

has a geometric distribution with $\Pr(J = j) = (1 - \rho_r)^{j-1} \rho_r$, where $j \geq 1$ and $\rho_r = 1 - (1 - \rho)^{\tilde{R}N}$. Consequently, the number of coded blocks per data packet possesses the memoryless property, a highly desirable attribute for the purpose of analysis.

This queueing system operates on top of the finite-state channel discussed earlier. In comparison to many previous studies, the resulting framework allows us to rigorously characterize the queueing performance while varying the block length and the code rate. The scaling property of the Poisson packet arrivals with geometric packet lengths is crucial in enabling the fair comparison of systems with different code parameters. In particular, the arrival process is defined at the bit level and we account for channel dependence within and across codewords. This last observation is especially pertinent for queueing systems, as correlation in service is known to exacerbate the distribution of a queue.

We emphasize that, in the current framework, a data packet is discarded from the transmit buffer if and only if the destination acknowledges reception of the latest codeword and this codeword contains the last parcel of information corresponding to the head packet. Packet departures are then determined by the channel realizations and the coding scheme. In particular, the code rate \tilde{R} has a major impact on performance. Generally, a lower code rate will have a small probability of decoding failure. However, it also needs more channel uses to complete the transmission of one data packet. Thus, for a fixed channel profile, we can vary the block length and code rate to find optimal system parameters. This natural tradeoff reflects the tension between the probability of a successful transmission and the size of its payload.

Let Q_s denote the number of data packets waiting in the transmitter queue after s codeword transmission intervals. The channel state at the same time instant is represented by C_{sN+1} . Notice that the channel state evolves more rapidly than events taking place in the queue. This explains the discrepancy between the indices. Based on these quantities, it is possible to define a Markov chain $U_s = (C_{sN+1}, Q_s)$ that captures the joint evolution of the queue and the channel over time. The ensuing transition probabilities from U_s to U_{s+1} are equal to

$$\Pr(U_{s+1} = (d, q_{s+1}) | U_s = (c, q_s)) = \sum_{n_1=0}^N P_{Q_{s+1}|N_1, Q_s}(q_{s+1}|n_1, q_s) P_{N_1, C_{(s+1)N+1}|C_{sN+1}}(n_1, d|c),$$

where the second term in the summand is given in Lemma 2. We can rewrite $P_{Q_{s+1}|N_1, Q_s}(q_{s+1}|n_1, q_s)$

as

$$\begin{aligned}
P_{Q_{s+1}|N_1, Q_s}(q_{s+1}|n_1, q_s) &= \sum_{e_1=0}^{n_1} \sum_{e_2=0}^{n_2} P_{Q_{s+1}, E_1, E_2|N_1, Q_s}(q_{s+1}, e_1, e_2|n_1, q_s) \\
&= \sum_{e_1=0}^{n_1} \sum_{e_2=0}^{n_2} P_{Q_{s+1}|E_1, E_2, N_1, Q_s}(q_{s+1}|e_1, e_2, n_1, q_s) P_{E_1, E_2|N_1, Q_s}(e_1, e_2|n_1, q_s) \\
&= \sum_{e_1=0}^{n_1} \sum_{e_2=0}^{n_2} \binom{n_1}{e_1} \binom{n_2}{e_2} \varepsilon_1^{e_1} (1 - \varepsilon_1)^{n_1 - e_1} \varepsilon_2^{e_2} (1 - \varepsilon_2)^{n_2 - e_2} P_{Q_{s+1}|E_1, E_2, N_1, Q_s}(q_{s+1}|e_1, e_2, n_1, q_s).
\end{aligned}$$

Suppose that the number of packets in the queue is $Q_s = q_s$, where $q_s > 0$. Then, admissible values for Q_{s+1} are restricted to the set $\{q_s - 1, q_s, q_s + 1, \dots\}$. The transition probabilities for $q_s > 0$ and $i \geq 0$ are given by

$$\begin{aligned}
P_{Q_{s+1}|E_1, E_2, N_1, Q_s}(q_s + i|e_1, e_2, n_1, q_s) \\
&= a_i P_{e|E_1, E_2, N_1}(e_1, e_2, n_1) + a_i (1 - P_{e|E_1, E_2, N_1}(e_1, e_2, n_1)) (1 - \rho_r) \\
&\quad + a_{i+1} (1 - P_{e|E_1, E_2, N_1}(e_1, e_2, n_1)) \rho_r
\end{aligned} \tag{30}$$

and the probability of the queue decreasing is

$$P_{Q_{s+1}|E_1, E_2, N_1, Q_s}(q_s - 1|e_1, e_2, n_1, q_s) = a_0 (1 - P_{e|E_1, E_2, N_1}(e_1, e_2, n_1)) \rho_r. \tag{31}$$

The queue can only become smaller when there are no arrivals, a codeword is successfully received at the destination, and the decoded codeword contains the last piece of data associated with a packet. Above, $P_{e|E_1, E_2, N_1}(e_1, e_2, n_1)$ is the conditional probability of decoding failure which appears in (24). The terms a_i denotes the probability that i packets arrive within the span of a codeword transmission. Since arrivals form a Poisson process, we have $a_i = \frac{(\lambda N)^i}{i!} e^{-\lambda N}$ for $i \geq 0$. When the queue is empty, $q_s = 0$, (30) applies for cases where $i \geq 1$. However, for $i = 0$, the conditional transition probability reduces to

$$P_{Q_{s+1}|E_1, E_2, N_1, Q_s}(0|e_1, e_2, n_1, 0) = a_0 + a_1 (1 - P_{e|E_1, E_2, N_1}(e_1, e_2, n_1)) \rho_r.$$

The overall profile of this system can be categorized as an M/G/1-type queue. The repetitive structure enables us to employ the matrix geometric method to compute the characteristics of this system and subsequently obtain its stationary distribution [28], [25].

VII. STOCHASTIC DOMINANCE

When the number of channel states is large, it may be impractical to employ exact probabilities of decoding failure. Even for memoryless channels, finding explicit expressions for different encoding/decoding schemes can be difficult. In the face of such a challenge, it is customary to turn to upper bounds on

the probabilities of decoding failure to provide performance guarantees. Furthermore, one can employ such upper bounds to assess the queueing performance of the system through stochastic dominance. We emphasize that this type of argument is independent of our proposed bounds. Rather, it applies to any upper bound on the probability of decoding failure.

The evolution of the queue length is governed by the Lindley equation,

$$Q_{s+1} = (Q_s + A_s - D_s)^+ \triangleq \max\{0, Q_s + A_s - D_s\} \quad (32)$$

where A_s is the number of arrivals that occurred during time interval s , and D_s is an indicator function for the potential completion of a packet transmission within the same time period. In this queueing model, the only inherent effect of replacing the probability of decoding failure by an upper bound is a potential reduction in the value of D_s . Using an upper bound on the failure probability naturally gives rise to a new random process \tilde{Q}_s defined by

$$\tilde{Q}_{s+1} = (\tilde{Q}_s + A_s - \tilde{D}_s)^+ \quad (33)$$

where \tilde{D}_s is drawn according to the distribution implied by the upper bound. We wish to show that, conditioned on starting in the same state and given a shared channel trace, the distribution of \tilde{Q}_s offers a conservative estimate of Q_s . To make this statement precise, we turn to an establish concept in probability.

Definition 1: A random variable Z is stochastically dominated by another random variable \tilde{Z} , a relation which we denote by $Z \preceq \tilde{Z}$, provided that

$$\Pr(Z > z) \leq \Pr(\tilde{Z} > z) \quad (34)$$

for all $z \in \mathbb{R}$. This relation extends to conditional probability laws. Suppose that

$$\Pr(Z > z|A) \leq \Pr(\tilde{Z} > z|A) \quad (35)$$

for all $z \in \mathbb{R}$. Then, we say that Z is stochastically dominated by \tilde{Z} , given A . We write this relation as $Z \preceq_A \tilde{Z}$.

A comprehensive discussion of stochastic dominance can be found in [29], [30]. In order to formulate the results we are interested in, we need to start by introducing two lemmas.

Lemma 4: The stochastic order defined in (34) is preserved under the positive part operation, $(\cdot)^+ = \max\{0, \cdot\}$.

Proof: From [29, Sec. 1.A.1], we know that given random variables X and Y , $X \preceq Y$ if and only if

$$\mathbb{E}[\varphi(X)] \leq \mathbb{E}[\varphi(Y)]$$

for all increasing functions $\varphi(\cdot)$ for which the expectations exists. In particular, $\varphi(\cdot) = \max\{0, \cdot\}$ is an increasing function. Thus, if $X \preceq Y$ then one can conclude that $X^+ \preceq Y^+$, as desired. ■

We explore the structure of potential packet completion events below. In establishing stochastic dominance, we will condition on a specific channel trace,

$$\vec{C} = \{C_1, C_{N+1}, C_{2N+1}, \dots\}$$

This ensures that the two stochastic processes experience a same level of difficulty at every step while trying to decode codewords. Note that, for our purposes, the dependence of Q_s and \tilde{Q}_s on \vec{C} is only through decoding attempts and, as such, this dependence is localized in time.

Lemma 5: Using upper bounds on the probabilities of decoding failure leads to stochastic dominance in potential packet completions, $\tilde{D}_s \preceq_{\vec{C}} D_s$.

Proof: A potential completion occurs when a codeword is decoded successfully at the destination, and the data it contains is the last segment of a packet. The probability of the latter event is ρ_r , and it is common to both D_s and \tilde{D}_s . However, the conditional probabilities of decoding failure differ, with $P_{e|\vec{C}=\vec{c}} \leq \tilde{P}_{e|\vec{C}=\vec{c}}$. This, in turn, gives

$$\Pr(\tilde{D}_s = 1 | \vec{C} = \vec{c}) = (1 - \tilde{P}_{e|\vec{C}=\vec{c}}) \rho_r \leq (1 - P_{e|\vec{C}=\vec{c}}) \rho_r = \Pr(D_s = 1 | \vec{C} = \vec{c}).$$

Since $D_s, \tilde{D}_s \in \{0, 1\}$, this equation is enough to establish stochastic ordering. ■

Collecting these results, we can turn to the behavior of the system. For a fair comparison, we assume that the two queues have the same number of packets at the onset of the communication process.

Proposition 6: Suppose that $Q_0 = \tilde{Q}_0$. Then, the process Q_s is stochastically dominated by \tilde{Q}_s .

Proof: As a first step, we assume that channel trace $\vec{C} = \vec{c}$ is fixed. We prove conditional dominance through mathematical induction. The base case follows from the condition of the theorem. As an inductive hypothesis, suppose that $Q_s \preceq_{\vec{C}} \tilde{Q}_s$. By Lemma 5, we have $\tilde{D}_s \preceq_{\vec{C}} D_s$. Since negation reverses stochastic dominance [29, pp. 9], we deduce that

$$-D_s \preceq_{\vec{C}} -\tilde{D}_s.$$

This, in turn, yields the relation

$$Q_s + A_s - D_s \preceq_{\vec{C}} \tilde{Q}_s + A_s - \tilde{D}_s.$$

The last step leverages the closure property of stochastic orders under convolutions [29, Thm. 1.A.3(b)].

Applying the increasing function $\varphi(\cdot) = \max\{0, \cdot\}$ to both sides, we immediately get

$$(Q_s + A_s - D_s)^+ \preceq_{\vec{C}} (\tilde{Q}_s + A_s - \tilde{D}_s)^+$$

from Lemma 4. That is, $Q_{s+1} \preceq_{\tilde{C}} \tilde{Q}_{s+1}$, thereby establishing our inductive step.

At this point, the statement of the proposition can be obtained by taking expectations over channel traces. Since the channel states are independent of queue sizes and code generation, the probabilistic weighing is the same for both Q_s and \tilde{Q}_s . This guarantees that the stochastic ordering is preserved. In other words, the relation $Q_s \preceq \tilde{Q}_s$ holds at all times. ■

There is a subtle distinction in the argument presented above. The random variable D_s indicates potential completion. Actual departures from the queue only take place when there are packets awaiting transmission. Mathematically, the distinction is resolved through the positive part operation. Conceptually, when the queue is empty, the source attempts to send a virtual packet with no physical meaning. This object is created for mathematical convenience.

In view of the aforementioned results, we can also consider the interplay between stochastic dominance and the stationary distributions of the queues. When the Markov chains U_s and \tilde{U}_s are positive recurrent, the corresponding queueing processes Q_s and \tilde{Q}_s are stable. In such cases, the stationary distribution associated with \tilde{Q}_s dominates the stationary distribution of Q_s . In particular, for any integer q , we have $\Pr(Q_s > q) \leq \Pr(\tilde{Q}_s > q)$ and, hence, in the limit we obtain

$$\begin{aligned} \Pr(Q > q) &= \lim_{s \rightarrow \infty} \Pr(Q_s > q) \\ &\leq \lim_{s \rightarrow \infty} \Pr(\tilde{Q}_s > q) = \Pr(\tilde{Q} > q) \end{aligned}$$

That is, $Q \preceq \tilde{Q}$ where Q and \tilde{Q} denote the stationary distributions of the two queueing processes listed above.

From a more intuitive point of view, one can argue that, pathwise, increasing the probability of failure can only result in fewer departures and, as such, there will remain at least as many packets waiting in the queue. In other words, when comparing two queueing systems with a same arrival process, a same underlying channel, and a same code generator, more decoding failures can only exacerbate the size of the queue. This observation holds in some generality and can be employed when the exact decoding error probability are not known or difficult to compute. This approach allows one to provide performance guarantees for a queueing system using bounds on the probabilities of decoding failure.

VIII. NUMERICAL RESULTS

In this section, we present numerical results for probabilities of decoding error and we compare them to the derived upper bounds. We also evaluate queueing performance using the exact error probabilities and their upper bounds derived in the rare-transition regime.




Fig. 2. Comparison of the approximate upper bound (14) with the exact bound (9) in the rare-transition regime with $N[\mathbf{P}_N]_{12} \approx 4$ and $N[\mathbf{P}_N]_{21} \approx 6$.

A. Comparison of Exponential Upper Bounds

We consider a communication system which transmits data over a Gilbert-Elliott channel with crossover probabilities $\varepsilon_1 = 0.01$ and $\varepsilon_2 = 0.1$. Figure 2 shows the approximate upper bounds of (14) as functions of block length and code rate, and compares them to the standard Gallager-type bounds of (9). Each curve shows the value of the bound averaged over all possible state transitions. Although the block lengths are relatively short, the approximate bounds are very close to the standard Gallager-type bounds. Furthermore, the difference becomes more negligible as N grows larger.

In Fig. 3, we plot the probabilities of decoding failure for the maximum-likelihood and minimum-distance decoders given by (19)-(23), against the bounds provided in (14). As anticipated, the maximum-likelihood decision rule outperforms the minimum distance decoder. For fixed N , there is a roughly constant ratio between the approximate upper bounds and the exact probabilities of error under maximum-likelihood decoding. This is not too surprising as similar statements can be made about the accuracy of Gallager-type bounds. We note that the figure features particularly short block lengths, as it is impractical to compute exact performance for long lengths.

B. Evaluation of Queueing Performance

We turn to the evaluation of overall performance and we consider a situation where, on average, packets are generated every 20 msec. This yields a rate of $\lambda = 50$ packets per second for the arrival process. The symbol rate for our binary channel is set to 28.75 Kb per second, which leads to an expected $1/575$




Fig. 3. Comparison of the approximate upper bound (14) with the exact probabilities of decoding failure under maximum-likelihood (ML) and minimum-distance (MD) decoding for $[\mathbf{P}_N]_{12} = 0.0533$ and $[\mathbf{P}_N]_{21} = 0.08$.

packets per channel use. The cross-over probabilities of the Gilbert-Elliott channel are set to $\varepsilon_1 = 0.01$ and $\varepsilon_2 = 0.1$, and its state transition probabilities are $\alpha = 0.0533$ and $\beta = 0.08$. Shannon capacity when the state is known at the receiver is therefore equal to 0.764 bits per channel use.

Increasing code rate \tilde{R} for a fixed block length decreases redundancy and therefore reduces the error-correcting capability of the code. Thus, the probability of decoding failure becomes larger. At the same time, changes in code rate affect ρ_r , the probability with which a codeword contains the last parcel of information of a packet. As code rate varies, these two phenomena alter the transition probabilities and, hence, they influence the stationary distribution of the Markov system in opposite ways.

The choice of a Poisson arrival process allows us to make fair comparisons between codes with different block lengths. In particular, the rate λ in packets per channel use is fixed, and arrivals in the queue correspond to the number of packets produced by the source during the transmission time of one codeword. The marginal distribution of the sampled process is Poisson with arrival rate λN , in packets per codeword. This formulation is new, and it bridges coding decision to queueing behavior in a rigorous manner. To examine overall system performance, we assume the existence of a genie which informs the receiver when an undetected decoding error occurs. Undetected errors are intrinsic to error channels, and this approach is standard when it comes to analysis. Still, for consistency, we require the system to feature a very low probability of undetected error, e.g., less than 10^{-5} , by a proper choice of the safety margin. That is, we only consider (N, \tilde{R}) pairs that meet this additional constraint.

Given this framework, a prime goal is to minimize the tail probability of the queue over all admissible

values of N , \tilde{R} , and τ or ν that satisfy the constraint on undetected error. To perform this task using the approximate exponential bound, we first evaluate the bound on undetected error probabilities for different rates and for $\tau = 0$ in (28). Then, for rates with high probability of undetected error, we increase τ so that the bound on probability of undetected error is decreased. Recall that this increases the probability of decoding failure. As we are also interested in minimizing the latter probability, we increase τ until the system meets the error-detecting condition and then stop. The values of N and \tilde{R} for which this procedure gives poor performance are ignored. A similar approach is used for system evaluation with exact error probabilities by changing the value of ν in (24), (27).

Figure 4 shows the approximate probability of the queue exceeding a threshold as a function of system parameters. The constraint on the number of packets in the queue is set to five, which reflects our emphasis on delay-sensitive communication. We have chosen τ in (28)–(29) such that $\max_{i,j} \tilde{P}_{\text{ue},S_N|S_0}(j|i)$ remains below 10^{-5} . The code rate considered vary from 0.25 to 0.75, with a step size of 0.05. Each curve corresponds to a different block length. As seen on the graph, there is a natural tradeoff between the probability of decoding failure and the payload per codeword. For a fixed block length, neither the smallest segment length nor the largest one delivers optimal performance. Moreover, block length must be selected carefully; longer codewords do not necessarily yield better queueing performance as they may result in large decoding delays. As such, the tail probability has a minimum over all rates and block lengths. Therefore, there are interior optimum points for both N and \tilde{R} . We see in Fig. 4 that the optimum code parameters are close to $(N, \tilde{R}) = (170, 0.5)$. For this particular set of code parameters, we have $\tau = 0.048$.

Figure 5 offers similar plots for the exact failure probability. Again, the optimum code parameters are near $(N, \tilde{R}) = (170, 0.5)$. In this case, $\nu = 8$ is the smallest value of ν that keeps $\max_{i,j} \bar{P}_{\text{ue},S_N|S_0}(j|i)$ below the 10^{-5} threshold. As we can see by comparing the results, performance evaluation based on the bound gives very good estimates for optimum coding parameters and overall system performance. Not only does the approximate bound give a good estimate of performance, it accurately predicts ideal system parameters for code block as small as 125. In addition, since the approximate bounds are slightly pessimistic, they produce conservative estimate of overall performance. Empirically, the systems perform better than predicted by the approximate error bounds.

IX. CONCLUSION

The rare-transition regime is a powerful methodology to characterize communication systems where the block length is of the same order or smaller than the coherence time of the channel. This mode




Fig. 4. Approximate bounds on the probability of the queue exceeding a threshold as functions of block length N and code rate \tilde{R} . The system parameters considered above are subject to $\max_{i,j} \tilde{P}_{ue,S_N|S_0}(j|i) \leq 10^{-5}$.




Fig. 5. Exact probability of the queue exceeding a threshold as functions of block length N and code rate \tilde{R} . The system parameters presented above meet the constraint $\max_{i,j} \tilde{P}_{ue,S_N|S_0}(j|i) \leq 10^{-5}$.

of operation is common in many practical implementations. This fact serves as a motivation for the proposed framework. In this article, we derived an approximate upper bound specifically tailored to the rare-transition regime to estimate the probability of decoding error over finite-state channels. A key property of the proposed methodology is that channel dependence within and across codewords is accounted for in the analysis. Furthermore, the proposed bound is numerically efficient to compute. It

can be employed for parameter selection and performance analysis in communication links subject to queueing constraints.

We provided supportive evidence for the accuracy of the bounding methodology by deriving exact expressions for the Gilbert-Elliott channel. Both maximum-likelihood decoding and a minimum-distance decision rule were considered. A numerical comparison between exact and approximate results validates our approach, showcasing the predictive power of the approximate bounding techniques. The numerical study focused on a two-state Markov channel formulation where state information is available at the destination. The methodology was subsequently extended to performance criteria based on the queueing behavior of the system. We provided a practical method to choose the block length and code rate as to minimize the probability that the transmit buffer exceeds a prescribed threshold. This is especially pertinent for communication links that support delay-sensitive traffic, yet it applies to general data stream as delay is known to negatively affect the performance of flow control and congestion control protocols. Numerical studies suggest that, for fixed conditions, optimal system parameters are essentially unaffected by small variations in the buffer overflow threshold.

The methodology and results are developed for finite-state Markov channels, but can be generalized to more intricate channels with memory, with or without symmetry property. In addition, the performance characterization of random codes over finite-state channels may extend to more practical schemes, such as iterative decoding of the Low Density Parity-Check (LDPC) codes. Possible avenues of future research further include the performance analysis of the rare-transition regime in the absence of side information at the receiver. Our conjecture is that in the limiting case, the channel sojourn time in each state is long enough for the receiver to estimate the state. For instance, for a channel with high correlation, patterns of errors with the same number of errors within a block are not equally likely. In fact, the system is more prone to burst of errors when the channel quality is poor. In other words, long channel memory enables the receiver to predict the channel quality. Hence, one might expect similar performance when the state information is not provided at the receiver in the rare-transition regime.

X. APPENDIX

A. Proof of Lemma 1

Distributions of the occupancy times for two-state discrete-time and continuous-time Markov chains have been studied previously. These distributions can be derived using bivariate generating functions and two-dimensional Laplace transforms, respectively [23]. Herein, we show how to adapt these approaches to derive the conditional distributions needed in our work.

The matrix of two-dimensional Laplace transforms for the distribution of the time spent in the first state over the time interval $[0, 1]$ is given by

$$\left[- \left(\mathbf{Q} - \begin{bmatrix} \theta & 0 \\ 0 & 0 \end{bmatrix} - \phi \mathbf{I} \right) \right]^{-1}.$$

For example, the first entry in the matrix is equal to

$$\frac{1}{u} + \frac{\mu \xi}{u(uv - \mu \xi)}$$

where $u = \phi + \theta + \mu$ and $v = \phi + \xi$. The inverse two-dimensional Laplace transform of this entry gives the conditional distribution $f_{\eta_1, S_f | S_i}(\cdot, 1|1)$. After this step, Lemma 2 in [23] can be employed to get the desired format in terms of modified Bessel functions. These are the expressions presented in Lemma 1.

B. Proof of Lemma 2

Let a and b be the numbers of transitions into and out of the initial state, respectively. Then, we can write $c = a + b$ to denote the total number of transitions that occur up to time N . From [23], we gather that

$$P_{N_1 | S_0}(m|1) = (1 - \alpha)^m (1 - \beta)^{N-m} \sum_{c=0}^{c_1} \binom{m}{a} \binom{N-m-1}{b-1} \left(\frac{\alpha}{1-\beta} \right)^b \left(\frac{\beta}{1-\alpha} \right)^a$$

where

$$c_1 = \begin{cases} N + \frac{1}{2} - |2m - \frac{1}{2} + N|, & m < N \\ 0, & m = N. \end{cases}$$

We can split the summation into two parts, one for odd and one for even values of c . If $c = 2k$, then $a = b = k$, and the corresponding sum represents $P_{N_1, S_N | S_0}(m, 1|1)$. If $c = 2k + 1$, then $a = k$, $b = k + 1$, and the resulting sum is $P_{N_1, S_N | S_0}(m, 2|1)$. As such, we can write

$$\begin{aligned} P_{N_1 | S_0}(m|1) &= (1 - \alpha)^m (1 - \beta)^{N-m} \\ &\times \left(\sum_{k=0}^{\infty} \binom{m}{k} \binom{N-m-1}{k-1} \left(\frac{\alpha}{1-\beta} \right)^k \left(\frac{\beta}{1-\alpha} \right)^k \right. \\ &\quad \left. + \sum_{k=0}^{\infty} \binom{m}{k} \binom{N-m-1}{k} \left(\frac{\alpha}{1-\beta} \right)^{k+1} \left(\frac{\beta}{1-\alpha} \right)^k \right). \end{aligned}$$

We can set the upper and lower limits on k to 0 and ∞ , respectively, since all other terms are automatically zero. From the definition of ${}_2F_1(-N + m + 1, -m; 1; \lambda)$ in [23, Lem. 1, pp. 383], we see

that

$$\begin{aligned} & \left(\frac{\alpha}{1-\beta} \right) \sum_{k=0}^{\infty} \binom{m}{k} \binom{N-m-1}{k} \left(\frac{\alpha}{1-\beta} \right)^k \left(\frac{\beta}{1-\alpha} \right)^k \\ &= \left(\frac{\alpha}{1-\beta} \right) {}_2F_1(-N+m+1, -m; 1; \lambda). \end{aligned}$$

Collecting these results, we obtain

$$\begin{aligned} P_{N_1, S_N | S_0}(m, 2|1) &= (1-\alpha)^m (1-\beta)^{N-m} \left(\frac{\alpha}{1-\beta} \right) \\ &\quad \times {}_2F_1(-N+m+1, -m; 1; \lambda) \end{aligned}$$

for $m = 1, 2, \dots, N-1$. Clearly, for $m = 0$ and $m = N$, this conditional probability is equal to zero.

Leveraging [23] and observing that

$$P_{N_1 | S_0}(m|1) = P_{N_1, S_N | S_0}(m, 1|1) + P_{N_1, S_N | S_0}(m, 2|1),$$

we can write the simplified equation

$$\begin{aligned} P_{N_1, S_N | S_0}(m, 1|1) &= (1-\alpha)^m (1-\beta)^{N-m} \\ &\quad \times ({}_2F_1(-N+m, -m; 1; \lambda) - {}_2F_1(-N+m+1, -m; 1; \lambda)) \end{aligned}$$

for $m = 1, 2, \dots, N-1$. Moreover, $P_{N_1, S_N | S_0}(0, 1|1) = 0$ and $P_{N_1, S_N | S_0}(N, 1|1) = (1-\alpha)^N$. The remaining conditional probabilities can be derived in the same manner.

REFERENCES

- [1] R. M. Fano, *Transmission of Information*. Cambridge, MA: The M.I.T. Press, 1961.
- [2] R. G. Gallager, *Information Theory and Reliable Communication*. New York, NY, USA: Wiley, 1968.
- [3] A. Barg and G. D. Forney, "Random codes: Minimum distances and error exponents," *IEEE Trans. Inform. Theory*, vol. 48, no. 9, pp. 2568–2573, 2002.
- [4] A. Sahai, "Why do block length and delay behave differently if feedback is present?" *IEEE Trans. Inform. Theory*, vol. 54, no. 5, pp. 1860–1886, 2008.
- [5] Y. Polyanskiy, H. V. Poor, and S. Verdú, "Channel coding rate in the finite blocklength regime," *IEEE Trans. Inform. Theory*, vol. 56, no. 5, pp. 2307–2359, 2010.
- [6] —, "Dispersion of the Gilbert-Elliott channel," *IEEE Trans. Inform. Theory*, vol. 57, no. 4, pp. 1829–1848, 2011.
- [7] A. Martinez, "Saddlepoint approximation of random-coding bounds," in *Proc. Annual Workshop on Inform. Theory and its Appl.*, 2011, pp. 1–6.
- [8] C. Nair, E. Ordentlich, and T. Weissman, "Asymptotic filtering and entropy rate of a hidden Markov process in the rare transitions regime," in *Proc. IEEE Int. Symp. Inform. Theory*, 2005, pp. 1838–1842.
- [9] Y. Peres and A. Quas, "Entropy rate for hidden Markov chains with rare transitions," *arXiv preprint arXiv:1012.2086*, 2010.

- [10] M. Asadi, R. P. Torghabeh, and N. P. Santhanam, "Estimation in slow mixing, long memory channels," *arXiv preprint arXiv:1301.6798*, 2013.
- [11] M. S. Pinsker, V. Prelov, and E. C. van der Meulen, "Transmission of a slowly varying Markov signal over memoryless channels," in *Proc. IEEE Int. Symp. Inform. Theory*, 2000, p. 488.
- [12] E. N. Gilbert, "Capacity of a burst-noise channel," *The Bell Syst. Techn. J.*, vol. 39, pp. 1253–1265, Sept. 1960.
- [13] E. O. Elliott, "Estimates of error rates for codes on burst-noise channels," *The Bell Syst. Techn. J.*, vol. 42, pp. 1977–1997, Sept. 1963.
- [14] J. R. Norris, *Markov chains*. Cambridge, MA: Cambridge University Press, 1998, no. 2008.
- [15] B. Sericola, "Occupation times in Markov processes," *Stochastic Models*, vol. 16, no. 5, pp. 479–510, 2000.
- [16] M. Bladt, B. Meini, M. F. Neuts, and B. Sericola, "Distributions of reward functions on continuous-time Markov chains," *Matrix-analytic methods: theory and application*, pp. 39–62, 2002.
- [17] W. Rudin, *Principles of Mathematical Analysis (International Series in Pure & Applied Mathematics)*. New York, NY: McGraw-Hill, 1976.
- [18] S. I. Resnick, *A probability path*. Boston, MA: Birkhäuser, 1999.
- [19] P. Billingsley, *Convergence of probability measures*. New York, New York, USA: John Wiley & Sons, 1968.
- [20] S. M. Ross, *Stochastic Processes*. New York, NY: John Wiley and Sons, 1983.
- [21] P. Bremaud, *Markov chains: Gibbs fields, Monte Carlo simulation, and queues*. New York, NY: Springer, 1999, vol. 31.
- [22] K. R. Gabriel, "The distribution of the number of successes in a sequence of dependent trials," *Biometrika*, vol. 46, no. 3, pp. 454–460, 1959.
- [23] P. J. Pedler, "Occupation times for two-state Markov chains," *Journal of Applied Probability*, pp. 381–390, 1971.
- [24] Y. Kovchegov, N. Meredith, and E. Nir, "Occupation times and bessel densities," *Statistics & Probability Letters*, vol. 80, no. 2, pp. 104–110, 2010.
- [25] F. Hamidi-Sepehr, J. F. Chamberland, and H. D. Pfister, "Delay-sensitive communication over fading channel: queueing behavior and code parameter selection," *arXiv preprint arXiv:1309.3307*, 2013.
- [26] D. Forney, "Exponential error bounds for erasure, list, and decision feedback schemes," *IEEE Trans. Inform. Theory*, vol. 14, no. 2, pp. 206–220, 1968.
- [27] E. Hof, I. Sason, and S. Shamai, "Performance bounds for erasure, list, and decision feedback schemes with linear block codes," *IEEE Trans. Inform. Theory*, vol. 56, no. 8, pp. 3754–3778, 2010.
- [28] A. Riska and E. Smirni, "Exact aggregate solutions for M/G/1-type Markov processes," in *Proc. of the ACM SIGMETRICS Int. Conf. on Meas. and Modeling of Comp. Syst.*, 2002, pp. 86–96.
- [29] M. Shaked and J. G. Shanthikumar, *Stochastic Orders and Their Applications*. New York, NY: Springer, 2006.
- [30] D. Stoyan, *Comparison methods for queues and other stochastic models*. New York, NY: John Wiley & Sons, 1983.